

II. Description of Program Activities

This section corresponds to the predefined forms required by the Division of Research Resources to provide information about our resource activities for their computerized retrieval system. These forms have been submitted separately and are not reproduced here to avoid redundancy with the more extensive narrative information about our resource and progress provided in this report.

II.A. Scientific Subprojects

Our core research and development activities are described starting on page 14, our training activities are summarized starting on page 97, and the progress of our collaborating projects is detailed starting on page 127.

II.B. Books, Papers, and Abstracts

The list of recent publications for our core research and development work starts on page 83 and those for the collaborating projects are in the individual reports starting on page 127.

II.C. Resource Summary Table

The details of resource usage, including a breakdown by the various subprojects, is given in the tables starting on page 100.

III. Narrative Description

III.A. Summary of Research Progress

III.A.1. Resource Overview

This is an annual report for year 15 of the SUMEX-AIM resource (grant RR-00785), the second year of a 5-year renewal period to support further research on applications of artificial intelligence in biomedicine. For the technical and administrative reasons discussed in earlier reports, the SUMEX project now includes the continuation of work on the development and dissemination of medical consultation systems (ONCOCIN) that had been supported before 1986 as resource-related research under grant RR-01631. Progress on core ONCOCIN research is therefore now reported here as well.

These combined efforts represent an ambitious research program to:

- Continue our long-range core research efforts on knowledge-based systems, aimed at developing new concepts and methodologies needed for biomedical applications.
- Substantially extend ONCOCIN research on developing and disseminating clinical decision support systems.
- Develop the core systems technology to move the national SUMEX-AIM community from a dependence on the central SUMEX DEC 2060 to a fully distributed, workstation-based computing environment.
- Introduce these systems technologies into the SUMEX-AIM community with appropriate communications and managerial assistance to responsibly phase out the central resource and DEC 2060 mainframe in a manner that will support community efforts to become self-sustaining and to continue scientific interactions through fully distributed means.
- Maintain our aggressive efforts at training and dissemination to help exploit the research potential of this field.

III.A.1.1. SUMEX-AIM as a Resource

SUMEX and the AIM Community

In the fifteen years since the SUMEX-AIM resource was established in late 1973, computing technology and biomedical artificial intelligence research have undergone a remarkable evolution and SUMEX has both influenced and responded to these changing technologies. It is widely recognized that our resource has fostered highly influential work in biomedical AI -- work from which much of the expert systems field emerged -- and that it has simultaneously helped define the technological base of applied AI research.

The focus of the SUMEX-AIM resource continues to emphasize research on artificial intelligence techniques that guide the design of computer programs that can help with the acquisition, representation, management, and utilization of the many forms of

medical knowledge in diverse biomedical research and clinical care settings -- ranging from biomolecular structure determination and analysis, to molecular biology, to clinical decision support, to medical education. Nevertheless, we have long recognized that the ultimate impact of this work in biomedicine will be realized through its assimilation with the full range of methodologies of medical informatics, such as data bases, biostatistics, human-computer interfaces, complex instrument control, and modeling. From the start, SUMEX-AIM work has been grounded in real-world applications, like systems for the interpretation of mass spectral information about biomolecular structures, chemical synthesis, interpretation of x-ray diffraction data on crystals, cognitive modeling, infectious disease diagnosis and therapy, DNA sequence analysis, experiment planning and interpretation in molecular biology, and medical instruction. Our current work extends this emphasis in application domains such as oncology protocol management, clinical decision support, protein structure analysis, and data base information retrieval and analysis. All of these research efforts have demanded close collaborations with diverse parts of the biomedical research community and the integration of many computational methods from those domains with knowledge-based approaches. Even though in the beginning the "AI-in-medicine" community was quite small, it is perforce no longer limited and easily-defined, but rather is spreading and is inextricably linked with the many biomedical applications communities we have collaborated with over the years. Driven both by the on-going diffusion of AI and by the development of personal computer workstations that signal the practical decentralization of computing resources, we must develop new resource communication and distributed computing technologies that will continue to facilitate wider intra- and inter-community communication, collaboration, and sharing of biomedical information.

The SUMEX Project has demonstrated that it is possible to operate a computing research resource with a national charter and that the services providable over networks were those that facilitate the growth of AI-in-Medicine. SUMEX now has a reputation as a model national resource, pulling together the best available interactive computing technology, software, and computer communications in the service of a national scientific community. Planning groups for national facilities in cognitive science, computer science, and biomathematical modeling have discussed and studied the SUMEX model and new resources, like the recently instituted BIONET resource for molecular biologists, are closely patterned after the SUMEX example.

The projects SUMEX supports have generally required substantial computing resources with excellent interaction. Today, with the dramatic explosion of high-performance workstations that are more and more generally available, the need for a central source of raw computing cycles has significantly diminished. In place of being a distributor of CPU cycles, SUMEX has become a communications cross-roads and a source of AI and computer systems software and expertise.

SUMEX has demonstrated that a computer resource is a useful "linking mechanism" for bringing together electronically teams of experts from different disciplines who share a common problem focus. AI concepts and software are among the most complex products of computer science. Historically it has not been easy for scientists in other fields to gain access to and mastery of them. Yet the collaborative outreach and dissemination efforts of SUMEX have been able to bridge the gap in numerous cases. About 40 biomedical AI application projects have developed in our national community and have been supported directly by SUMEX computing resources over the years -- many more have benefitted indirectly through access to the software, information, and advice offered by the SUMEX resource.

The integration of AI ideas with other parts of medical informatics and their dissemination into biomedicine is happening largely because of the development in the 1970's and early 1980's of methods and tools for the application of AI concepts

to difficult professional-level problem solving. Their impact was heightened because of the demonstration in various areas of medicine and other life sciences that these methods and tools really work. Here SUMEX has played a key role, so much so that it is regarded as "the home of applied AI."

SUMEX has been the home of such well-known AI systems as DENDRAL (chemical structure elucidation), MYCIN (infectious disease diagnosis and therapy), INTERNIST (differential diagnosis), ACT (human memory organization), MOLGEN/BIONET (tools for DNA sequence analysis and molecular biology experiment planning), ONCOCIN (cancer chemotherapy protocol advice), SECS (chemical synthesis), EMYCIN (rule-based expert system tool), and AGE (blackboard-based expert system tool). In the past four years, our community has published a dozen books that give a scholarly perspective on the scientific experiments we have been performing. These volumes, and other work done at SUMEX, have played a seminal role in structuring modern AI paradigms and methodology.

The Future of SUMEX-AIM

Given this background, what is the future need and course for SUMEX as a resource -- especially in view of the on-going revolution in computer technology and costs and the emergence of powerful single-user workstations and local area networking? The answers remain clear.

Basic Research on AI in Biomedicine

At the deepest research level, despite our considerable success in working on medical and biological applications, the problems we can attack are still sharply limited. Our current ideas fall short in many ways against today's important health care and biomedical research problems brought on by the explosion in medical knowledge and for which AI should be of assistance. Just as the research work of the 70's and 80's in the SUMEX-AIM community fuels the current practical and commercial applications, our work of the late 80's will be the basis for the next decade's systems.

The report of the panel on medical informatics [3], convened late in 1985 by the National Library of Medicine to review and recommend twenty-year goals for the NLM, listed among its highest priority recommendations the need to greatly expand and aggressively pursue an interdisciplinary research program to develop computational methods for acquiring, representing, managing, and using biomedical knowledge of all sorts for health care and biomedical research. These are precisely the problems which the SUMEX-AIM community has been working on so successfully and which will require work well beyond the five year funding period we have requested. It is essential that this line of research in the SUMEX-AIM community, represented by our core AI research, the ONCOCIN research, and our collaborative research groups, be continued.

The Changing Role of the Central Resource

At the resource level, there are changing, but still growing, needs for computing resources for the active AIM research community to continue its work over the next five years. The workstations to which we directed our attention in 1980 have now demonstrated their practicality as research tools and, increasingly, as potential mechanisms for disseminating AI systems as cost-effective decision aids in clinical settings such as private offices. The era of highly centralized general machines for AI research is rapidly coming to an end and is being replaced by networks of distributed but heterogeneous single-user machines sharing common information

resources and communication paths among members of the biomedical research community.

Most of our community groups have been able to take advantage of local computing facilities, with SUMEX-AIM providing a central cross-roads for communications and the sharing of programs and knowledge. In its core research and development role, SUMEX-AIM has its sights set on the hardware and software systems of the next decade. We expect major changes in the distributed computing environments that are just now emerging in order to make effective use of their power and to adapt them to the development and dissemination of biomedical AI systems for professional user communities. In its training role, SUMEX is a crucial resource for the education of badly needed new researchers and professionals to continue the development of the biomedical AI field. The "critical mass" of the existing physical SUMEX resource, its development staff, and its intellectual ties with the Stanford Knowledge Systems Laboratory (KSL -- see Appendix A for a summary of current KSL research activities), make this an ideal setting to integrate, experiment with, and export these methodologies for the rest of the AIM community.

We will continue our *experimental* approach to distributed systems, learning to build and exploit distributed networks of these machines and to build and manage graceful software for these systems. Since decentralization is central to our future, we must learn its technical characteristics.

Resource Sharing

An equally important function of the SUMEX-AIM resource is an exploration of the use of computer communications as a means for interactions and sharing between geographically remote research groups engaged in biomedical computer science research and for the dissemination of AI technology. This facet of scientific interaction is becoming increasingly important with the explosion of complex information sources and the regional specialization of groups and facilities that might be shared by remote researchers [2, 1]. Another of the key recommendations of the NLM medical informatics planning panel [3] was that high-speed network communication links be established throughout the biomedical research community so that knowledge and information can be shared across diverse research groups and that the required interdisciplinary collaborations can take place. Recent efforts to establish a national NSFNet, largely to support the supercomputer projects funded by NSF but also to replace and upgrade part of the national research community linkage that the now aging ARPANET has supported, have made important progress. Still, these efforts do not encompass the broad range of biomedical research groups that need national network access and to date, the NIH has not played an aggressive role in the interagency Research Internet coordination efforts. We must work to build a stronger institutional support for a National Research Network.

SUMEX continues to be an important pathfinder to develop the technology and community interaction tools needed to expand community system and communication resources. Our community building effort is based upon the developing state of distributed computing and communications technology and we have therefore turned our core systems research to actively supporting the development of distributed computing and communications resources to facilitate collaborative project research and continued inter-group communications.

Summary of Long-term Goals

- Maintain the synergistic relationship between SUMEX core system development, core AI research, our experimental efforts at disseminating clinical decision-making aids, and new applications efforts.
- Continue to serve the national AIM research community, less and less as a source of raw computing cycles and more and more as a transfer point for new technologies important for community research and communication. We will also continue our coordinating role within the community through electronic media and periodic AIM workshops.
- Maintain our connections to ARPANET, TELENET, and our local Ethernet and assist other community members to establish similar links by example, by integrating and providing enabling software, and by offering advice and support within our resources.
- Focus new computing resource developments on more effective exploitation of distributed workstations through better communication and cooperative computing tools, using transparent digital networking schemes.
- Enhance the computing environments of workstations so that only minimal dependency on central, general-purpose computing hosts remains and these mainframe time-sharing systems can be phased out eventually. Remaining central resources will include servers for communications, community information resources, and special computing architectures (e.g., shared- or distributed-memory symbolic multiprocessors) justified by cost-effectiveness and unique functionality.
- Incrementally phase in, disseminate, and evaluate those aspects of the local distributed computing resource that are necessary for continuing national AIM community support within this distributed paradigm. This will ultimately point the way towards the distributed computing resource model that we believe will interlink this community well into the next decade.
- Gradually and responsibly phase out the existing DEC 2060 machine as effective distributed computing alternatives become widely available. We expect this to be possible sometime during the next year of the continuation resource.
- Continue the central staff and management structure, essentially unchanged in size and function during the five-year transition period, except for the merging of the core part of the ONCOCIN research with the SUMEX resource.

III.A.1.2. Significance and Impact in Biomedicine

Artificial intelligence is the computer science of representations of symbolic knowledge and its use in symbolic inference and problem-solving processes. Projects in the SUMEX-AIM community are concerned in some way with the application of AI to biomedical research and the resource has given strong impetus and support to knowledge-based system research in biomedicine. For computer applications in medicine and biology, this research path is crucial. Medicine and biology are not presently mathematically-based sciences; unlike physics and

engineering, they are seldom capable of exploiting the mathematical characteristics of computation. They are essentially inferential, not calculational, sciences. If the computer revolution is to affect biomedical scientists, computers will be used as inferential aids.

The growth in medical knowledge has far surpassed the ability of a single practitioner to master it all, and the computer's superior information processing capacity thereby offers a natural appeal. Furthermore, the reasoning processes of medical experts are poorly understood; attempts to model expert decision-making necessarily require a degree of introspection and a structured experimentation that may, in turn, improve the quality of the physician's own clinical decisions, making them more reproducible and defensible. New insights that result may also allow us more adequately to teach medical students and house staff the techniques for reaching good decisions, rather than merely to offer a collection of facts which they must independently learn to utilize coherently.

Perhaps the larger impact on medicine and biology will be the exposure and refinement of the hitherto largely private heuristic knowledge of the experts of the various fields studied. The ethic of science that calls for the public exposure and criticism of knowledge has traditionally been flawed for want of a methodology to evoke and give form to the heuristic knowledge of scientists. AI methodology is beginning to fill that need. Heuristic knowledge can be elicited, studied, critiqued by peers, and taught to students.

The importance of AI research and its applications is increasing in general, without regard for the specific areas of biomedical interest. AI is one of the principal fronts along which university computer science groups are expanding. The pressure from student career-line choices is great. Federal and industrial support for AI research and applications is vigorous, although support specifically for biomedical applications continues to be limited. All of the major computer manufacturers (e.g., IBM, DEC, TI, UNISYS, HP, and others) are using and marketing AI technology aggressively and many software companies are putting more and more products on the market. Many other parts of industry are also actively pursuing AI applications in their own contexts, including defense and aerospace companies, manufacturing companies, financial companies, and others.

Despite the limited research funding available, there is also an explosion of interest in medical AI. The American Association for Artificial Intelligence (AAAI), the principal scientific membership organization for the AI field, has 7000 members, several thousand of whom are members of the medical special interest group known as the AAAI-M. Speakers on medical AI are prominently featured at professional medical meetings, such as the American College of Pathology and American College of Physicians meetings; a decade ago, the words *artificial intelligence* were never heard at such conferences. And at medical computing meetings, such as the annual Symposium on Computer Applications in Medical Care (SCAMC) and the international MEDINFO conferences, the growing interest in AI and the rapid increase in papers on AI and expert systems are further testimony to the impact that the field is having.

AI is beginning to have a similar effect on medical education. Such diverse organizations as the National Library of Medicine, the American College of Physicians, the Association of American Medical Colleges, and the Medical Library Association have all called for sweeping changes in medical education, increased educational use of computing technology, enhanced research in medical computer science, and career development for people working at the interface between medicine and computing. They all cite evolving computing technology and (SUMEX-AIM) AI research as key motivators. At Stanford, we have vigorous special programs for student training and research in AI -- a graduate program in Medical Information

Sciences and the two-year Masters Degree in AI program. All of these have many more applicants than available slots. Demand for their graduates, in both academic and industrial settings, is so high that students typically begin to receive solicitations one or two years before completing their degrees.

III.A.1.3. Summary of Current Resource Goals

The following outlines the specific objectives of the SUMEX-AIM resource during the current three-year award period begun in August 1986. It provides an overall research plan for the resource and provides the backdrop against which specific progress is reported. Note that these objectives cover only the resource nucleus; objectives for individual collaborating projects are discussed in their respective reports in Section IV. Specific aims are broken into five categories: 1) Technological Research and Development, 2) Collaborative Research, 3) Service and Resource Operations, 4) Training and Education, and 5) Dissemination.

1) Technological Research and Development

SUMEX funding and computational support for core research is complementary to similar funding from other agencies (including DARPA, NASA, NSF, NLM, private foundations, and industry) and contributes to the long-standing interdisciplinary effort at Stanford in basic AI research and expert system design. We expect this work to provide the underpinnings for increasingly effective consultative programs in medicine and for more practical adaptations of this work within emerging microelectronic technologies. Specific aims include:

- Basic research on AI techniques applicable to biomedical problems. Over the next term we will emphasize work on blackboard problem-solving frameworks and architectures, knowledge acquisition or learning, constraint satisfaction, and qualitative simulation.
- Investigate methodologies for disseminating application systems such as clinical decision-making advisors into user groups. This will include generalized systems for acquiring, representing and reasoning about complex treatment protocols such as are used in cancer chemotherapy and which might be used for clinical trials.
- Support community efforts to organize and generalize AI tools and architectures that have been developed in the context of individual application projects. This will include retrospective evaluations of systems like the AGE blackboard experiment and work on new systems such as BB1, CARE, EONCOCIN, EOPAL, Meta-ONYX, and architectures for concurrent symbolic computing. The objective is to evolve a body of software tools that can be used to more efficaciously build future knowledge-based systems and explore other biomedical AI applications.
- Develop more effective workstation systems to serve as the basis for research, biomedical application development, and dissemination. We seek to coordinate basic research, application work, and system development so that the AI software we develop for the next 5-10 years will be appropriate to the hardware and system software environments we expect to be practical by then. Our purchases of new hardware will be limited to experimentation with state-of-the-art workstations as they become available for our system developments.

2) Collaborative Research

- Encourage the exploration of new applications of AI to biomedical research and improve mechanisms for inter- and intra-group collaborations and communications. While AI is our defining theme, we may consider exceptional applications justified by some other unique feature of SUMEX-AIM essential for important biomedical research. We will continue to exploit community expertise and sharing in software development.
- Minimize administrative barriers to the community-oriented goals of SUMEX-AIM and direct our resources toward purely scientific goals. We will retain the current user funding arrangements for projects working on SUMEX facilities. User projects will fund their own manpower and local needs; actively contribute their special expertise to the SUMEX-AIM community; and receive an allocation of system resources under the control of the AIM management committees. We will begin charging "fees for service" to Stanford users as DRR support for the DEC 2060 is phased out. Fees to national users will be delayed as long as financially possible.
- Provide effective and geographically accessible communication facilities to the SUMEX-AIM community for remote collaborations, communications among distributed computing nodes, and experimental testing of AI programs. We will retain the current ARPANET and TELENET connections for at least the near term and will actively explore other advantageous connections to new communications networks and to dedicated links.

3) Service and Resource Operations

SUMEX-AIM does not have the computing or manpower capacity to provide routine service to the large community of mature projects that has developed over the years. Rather, their computing needs are better met by the appropriate development of their own computing resources when justified. Thus, SUMEX-AIM has the primary focus of assisting new start-up or pilot projects in biomedical AI applications in addition to its core research in the setting of a sizable number of collaborative projects. We do offer continuing support for projects through the lengthy process of obtaining funding to establish their own computing base.

4) Training and Education

- Provide documentation and assistance to interface users to resource facilities and systems.
- Exploit particular areas of expertise within the community for assisting in the development of pilot efforts in new application areas.
- Accept visitors in Stanford research groups within limits of manpower, space, and computing resources.
- Support the Medical Information Science and MS/AI student programs at Stanford to increase the number of research personnel available to work on biomedical AI applications.
- Support workshop activities including collaboration with other community

groups on the AIM community workshop and with individual projects for more specialized workshops covering specific research, application, or system dissemination topics.

5) *Dissemination*

While collaborating projects are responsible for the development and dissemination of their own AI systems and results, the SUMEX resource will work to provide community-wide support for dissemination efforts in areas such as:

- Encourage, contribute to, and support the on-going export of software systems and tools within the AIM community and for commercial development.
- Assist in the production of video tapes and films depicting aspects of AIM community research.
- Promote the publication of books, review papers, and basic research articles on all aspects of SUMEX-AIM research.

III.A.2. Details of Technical Progress

This section gives an overview of progress for the nucleus of the SUMEX-AIM resource. A more detailed discussion of our progress in specific areas and related plans for further work are presented beginning on Page 19. Objectives and progress for individual collaborating projects are discussed in their respective reports on Page 127. These collaborative projects collectively provide much of the scientific basis for SUMEX as a resource and our role in assisting them has been a continuation of that evolved in the past. Collaborating projects are autonomous in their management and provide their own manpower and expertise for the development and dissemination of their AI programs.

III.A.2.1. Progress Highlights

In this section we summarize highlights of SUMEX-AIM resource activities over the past year (May 1987 - April 1988), focusing on the resource nucleus.

- We made excellent progress in **core AI research**. We have begun to explore the design and use of very large knowledge bases with the hypothesis that both the problems of brittleness and over-specialization in current knowledge-based systems can be addressed by constructing large, multi-use knowledge bases (LMKB). A LMKB would 1) encode domain knowledge in greater depth and breadth than required for any specific task, 2) encode knowledge that cuts across many domains of expertise, and 3) serve as a core repository of knowledge to be accessed by large numbers of specific applications.

Research has also progressed on several fundamental issues of AI, including knowledge representation, blackboard frameworks, parallel symbolic computing architectures, and machine learning. Work continues in PROTEAN and BB1 on the explicit representation of control knowledge, on the representation of geometric problem-solving knowledge in PROTEAN and PEAKS, on the representation of diagnosis expertise and the integration of a numerical simulator of a model system with an expert system in ABLE, and on the flexible, rich representation of control knowledge to facilitate modeling of problem-solving at the strategic level as well as at the tactical level. We have continued to develop the BB1 blackboard architecture for systems that reason about (control, explain, and learn about) their own actions. The BB1 system is being applied in a new domain, the real time monitoring of patients in an intensive care unit (BBICU).

The parallel architectures work has developed a CAD (Computer Aided Design) system for hierarchical, multiple-level specification of computer architectures (SIMPLE); a parameterized, multiprocessor array emulation defined in SIMPLE's specification languages and running on SIMPLE's simulator (CARE); a set of extensions to Lisp for studying expressed concurrency in functional programming, object-oriented, and shared-variable models of concurrent computation (LAMINA); and two alternative parallel blackboard frameworks for expressing application problems (POLIGON and CAGE). These have been applied to several signal understanding problems with promising problem-solving speed-up.

The machine learning work has concentrated on explanation-based generalization and chunking work in the SOAR framework, inductive rule

learning, "apprenticeship" learning, and tools for acquiring knowledge and debugging knowledge structures. Our research in machine learning has focused on several distinct problem domains including medicine (NEOMYCIN/HERACLES), physics (ABLE), and biochemistry (PROTEAN) in addition to domain-independent investigations.

Work has also continued on reasoning with uncertainty to find ways of combining formal and informal approximate reasoning methods. Specifically, this research seeks to (1) develop techniques for using knowledge about problem-solving tradeoffs to dynamically optimize system performance for the user, (2) construct efficient algorithms for probabilistic reasoning, and (3) investigate pragmatic techniques for the elicitation of knowledge from experts.

- We have made continued significant progress in the **core ONCOCIN research** to generalize the tools for clinical trial management from the initial cancer chemotherapy management application. A major accomplishment this year from our work on the examination of the structures of protocols across medical subspecialties other than cancer chemotherapy (e.g., hypertension and insulin diabetes treatment) was the creation of a generalized knowledge acquisition tool designed to encode descriptions of clinical trials. The system is called PROTEGE and produces as its output a *computer program* which is an OPAL-like clinical trial protocol definition system specifically tailored for an arbitrary clinical area such as hypertension. This OPAL-like system can then be used to create the knowledge base (e.g., for hypertension) that drives an ONCOCIN-like patient/protocol management system.

We continued development of the OPAL system for graphical knowledge acquisition to facilitate protocol definition and knowledge base entry for the ONCOCIN oncology application area. We began to explore alternative platforms for developing OPAL-like systems such as HyperCard on the MAC II. We also tested the usefulness of a "generic protocol viewing tool" based on a relational data base storage mechanism. We performed several experiments aimed at developing a single underlying storage mechanism (termed a "cell") from which various interface systems including the Interviewer flowsheet, a generalized spreadsheet utility, and the OPAL schema entry flowchart interface could be derived.

We have begun a project to explore the integration of speech-recognition technology into the interface to ONCOCIN. The project uses a commercially available continuous speech recognition product and a prototype ONCOCIN adaptation now permits users to navigate the graphical interface and enter clinical data using spoken commands. The development of this system exploits the significant experience we have gained in distributed computing since the phonetic device, initial parsing software, and the ONCOCIN system all reside on different pieces of hardware.

We released a new version of our ONCOCIN object language which has proven to be the most stable and powerful version to date. We have also continued to examine the issues of disseminating the ONCOCIN system into actual clinical settings.

- We have made excellent progress on the **core system development** work targeted at supporting the distributed AIM community. We elaborated

our systems development plans further at a site visit held in August 1987 in response to our request to the National Advisory Research Resources Council to restore the final 2 years of our grant award. Following a special study section review and reconsideration by the Council, our plan and the full 5-year grant award were approved. In line with the review group recommendations, we have moved to sharply focus our development resources on a small number of standardized hardware and software configurations.

After consideration of many requirements and alternative systems for AIM community computing needs for the coming years to replace and upgrade the services of the 2060, we have chosen Apple Macintosh II workstations as the general computing environment for researchers and staff, TI Explorer Lisp machines (including the microExplorer Macintosh coprocessor) as the near-term high-performance Lisp research environment, and a SUN-4 as the central system network server (wide- and local-area network interfaces, file services, printing services, etc.). The bulk of this hardware was purchased with DARPA research funds and we are now beginning the installation and integration process. We have concentrated our current systems efforts on getting basic capabilities operational, such as for text processing, filing/archiving, printing, graphics, office management, system building tools, information resource access, and distributed system operation and management) etc.

Development work for the Mac/Explorer/SUN environments has been limited because of manpower cuts necessitated by NIH cuts in our funding awards. Available resources have focussed on providing remote access between workstations, integrating a solid support of the TCP-IP network protocol, and building a distributed electronic mail system. We have significantly refined the prototype distributed EMail system developed for the Xerox D-machines last year (a number of people are using the system routinely to manage their mail) and are porting this to the Explorer and adapting it for the Mac II.

We have started a project based on the MacWorkstation software licensed from APPLE which is designed to allow the Macintosh to start up programs on mainframes or other workstations and receive graphic and text output to the MAC in a seamless fashion. We plan to write a Common Lisp window package (based on Common Windows) that uses MacWorkstation so we can connect (via Ethernet, AppleNet or RS-232C/modem) to any of our Common Lisp engines and run the same piece of code on any of them.

One of the key issues in selecting the systems for our distributed computing environment was the performance of Common Lisp and to help make this evaluation, we undertook an informal survey of the performance of two KSL AI software packages, SOAR and BB1, on a wide variety of machines. Within a factor of two of the best performance, a considerable range of workstations based on stock microprocessor chips as well as specially microprogrammed Lisp chips have comparable performance. Even though performance gaps between microprogrammed Lisp systems and stock workstation implementations are narrowing, there still remains a significant difference in the quality of the development environments. We have attempted to distill the key features of the Lisp machine environments that would be needed in stock machine implementations in order to make them attractive in a development setting.

We have installed an interface between our Ethernet environment and the TELENET network by which many AIM users gain access to SUMEX that allows full access to distributed SUMEX resources. We have also continued to maintain the DEC 2060 environment to stay current with operating system and other environmental upgrades.

- We have continued the **dissemination of SUMEX-AIM technology** through various media. We have reorganized the distribution system for our AI software tools (EMYCIN, AGE, and BB1) to academic, industrial, and federal research laboratories, in order to make it more efficient and require less research staff time. We have also continued to distribute the video tapes of some of our research projects including ONCOCIN, and an overview tape of Knowledge Systems Laboratory work to outside groups. Our group has continued to publish actively on the results of our research, including more than 45 research papers per year in the AI literature and a dozen books in the past 5 years on various aspects of SUMEX-AIM AI research. We assisted and participated actively in the AIM Workshop sponsored by AAAI and held at Stanford (Ramesh Patil from MIT was the Program Chairman) and hosted a number of AIM community visitors at our Stanford research laboratory this past year.
- The **Medical Information Sciences program**, begun at Stanford in 1983 under Professor Shortliffe as Director, has continued its strong development over the past year. The specialized curriculum offered by the MIS program focuses on the development of a new generation of researchers able to support the development of improved computer-based solutions to biomedical needs. The feasibility of this program resulted in large part from the prior work and research computing environment provided by the SUMEX-AIM resource. It has recently received enthusiastic endorsement from the Stanford Faculty Senate for an additional five years, has been awarded renewed post-doctoral training support from the National Library of Medicine with high praise for the training and contributions of the SUMEX-AIM environment from the reviewing study section, and has received additional industrial and foundation grants for student support. This past year, MIS students have published many papers, including several that have won conference awards.
- We have continued to recruit **new user projects and collaborators** to explore further biomedical areas for applying AI. A number of these projects are built around the communications network facilities we have assembled, bringing together medical and computer science collaborators from remote institutions and making their research programs available to still other remote users. At the same time we have encouraged older mature projects to build their own computing environments thereby facilitating the transition to a distributed AIM community. A substantial number of projects have moved to their own computing resources, including SOAR, under Dr. Paul Rosenbloom at USC/ISI in Los Angeles; the Logic Group projects (DART, Intelligent Agents, and MRS) under Professor Michael Genesereth at Stanford; Hierarchical Models of Human Cognition (CLIPR), under Professors Walter Kintsch and Peter Polson at the University of Colorado; Problem Solving Expertise (SOLVER), under Professors Paul Johnson and William Thompson at the University of Minnesota; RXDX, under Professor Robert Lindsay at the University of Michigan; and Computer-Based Exercises in Pathophysiologic Diagnosis, under Professor J. Robert Beck at Dartmouth College.

SUMEX user projects have made good progress in developing and disseminating effective consultative computer programs for biomedical research. These systems provide expertise in areas like cancer chemotherapy protocol management, clinical diagnosis and decision-making, and molecular biology. We have worked hard to meet their needs and are grateful for their expressed appreciation (see Section IV).

III.A.2.2. Core ONCOCIN Research

ONCOCIN is a data management and therapy advising program for complex cancer chemotherapy experiments. The development of the system began in 1979, following the successful generalization of MYCIN into the EMYCIN expert system shell. The ONCOCIN project has evolved over the last eight years: the original version of ONCOCIN ran on the time-shared DEC computers using a standard terminal for the time-oriented display of patient data. The current version uses compact workstations running on the Ethernet network with a large bit-mapped displays for presentation of patient data. The project has also expanded in scope. There are three major research components: 1) ONCOCIN, the therapy planning program and its graphical interface; 2) OPAL, a graphical knowledge entry system for ONCOCIN; and 3) ONYX, a strategic planning program designed to give advice in complex therapy situations. Each of these research components has been split into two parts: continued development of the cancer therapy versions of the system, and generalization of each of the components for use in other areas of medicine. This annual report will describe our work on each of the components: implementation of ONCOCIN workstations in the Stanford clinic, knowledge acquisition research, and research to generalize ONCOCIN for application in clinical trial domains other than medical oncology (E-ONCOCIN).

A major highlight of this year was the creation of a generalized knowledge acquisition tool designed to encode descriptions of clinical trials. The system, named PROTEGE, was the Ph.D. thesis work of Mark Musen. The output of PROTEGE is an OPAL-like input system designed for one clinical area such as hypertension. This input system (HTN-OPAL) can then be used to create the hypertension knowledge base for an E-ONCOCIN like system. This experiment was carried out this year for both the hypertension and oncology domains. Details of this project are described later in this report.

1 - Overview of the ONCOCIN Therapy Planning System

ONCOCIN is an advanced expert system for clinical oncology. It is designed for use after a diagnosis has been reached, focusing instead on assisting with the management of cancer patients who are receiving chemotherapy. Because anticancer agents tend to be highly toxic, and because their tumor-killing effects are routinely accompanied by damage to normal cells, the rules for monitoring and adjusting treatment in response to a given patient's course over time tend to be complex and difficult to memorize. ONCOCIN integrates a temporal record of a patient's ongoing treatment with an underlying knowledge base of treatment protocols and rules for adjusting dosage, delaying treatment, aborting cycles, ordering special tests, and similar management details. The program uses such knowledge to help physicians with decisions regarding the management of specific patients.

A major lesson of past work in clinical computing has been the need to develop methods for integrating a system smoothly into the patient-care environment for which it is intended. In the case of ONCOCIN, the goal has been to provide expert consultative advice as a by-product of the patient data management process, thereby avoiding the need for physicians to go out of their way to obtain advice. It is intended that oncologists use ONCOCIN routinely for recording and reviewing patient data on the computer's screen, regardless of whether they feel they need decision-making assistance. This process replaces the conventional recording of data on a paper flowsheet and thus seeks to avoid being perceived as an additive task. In accordance with its knowledge of the patient's chemotherapy protocol, ONCOCIN then provides assistance by suggesting appropriate therapy at the time that the day's treatment is to be recorded on the flowsheet. Physicians maintain control of the

decision, however, and can override the computer's recommendation if they wish. ONCOCIN also indicates the appropriate interval until the patient's next treatment and reminds the physician of radiologic and laboratory studies required by the treatment protocol.

2 - Implementation of the ONCOCIN Workstation in the Stanford Clinic

In mid-1986, we placed the workstation version of ONCOCIN into the Oncology Day Care clinic. This version is a completely different program from the version of ONCOCIN that was available in the clinic from 1981-1985-- using protocols entered through the OPAL program, with a new graphical data entry interface, and revised knowledge representation and reasoning component. One person in the clinic (Andy Zelenetz) became primarily responsible for making sure that our design goals for this version of ONCOCIN were met. His suggestions included the addition of key protocols and the ability to have the program be useful for clinicians as a data management tool if the complete treatment protocol had not yet been entered into the system. Additional fellows are being trained on a very stable release of ONCOCIN that became available in early 1988. A version of the system was sent to the University of Pittsburgh for evaluation. In addition, ONCOCIN will become available for presentation at the National Library of Medicine Artificial Intelligence Demonstration Center. For these various efforts, Janice Rohn has created an extensive user manual, sample patient interactions, and reminder cards to shorten the training period for ONCOCIN.

The process of entering a large number of treatment protocols in a short period of time led to other research topics including: design of an automated system for producing meaningful test cases for each knowledge, modification of the design of the time-oriented database and the methods for accessing the database, and the development of methods for graphically viewing multiple protocols that are combined into one large knowledge base. These research efforts will continue into the next year. In addition, some of the treatment regimens developed for the original mainframe version are still in use and can be transferred to the new version of ONCOCIN.

We also received new insights about the design of the internal structures of the knowledge base (e.g., the relationship between the way we refer to chemotherapies, drugs, and treatment visits). We will continue to optimize the question-asking procedure, the method for traversing the plan structure in the knowledge base, and consider alternative arrangements used to represent the structure of chemotherapy plans. Although we have concentrated our review of the ONCOCIN design primarily on the data provided by additional protocols, we know that non-cancer therapy problems raise similar issues. The E-ONCOCIN effort is designed to produce a domain-independent therapy planning system that includes the lessons learned from our oncology research.

3 - E-ONCOCIN: Domain Independent Therapy Planning

During the past two years, our E-ONCOCIN research has concentrated on understanding how protocols in medicine vary across subspecialties. We are examining several application areas: the intensive care unit, insulin treatment for diabetes, hypertension protocols, and both standard and complex cancer treatment problems. The diagnosis and therapy selection for patients in the intensive care unit is a natural application area because it is based on changing data and the need to determine the response to therapy interventions. In addition, it is an area where reasonable mathematical models of the respiratory system can be integrated into the expert system. We also felt that the area of insulin treatment for diabetes would be a

good area to explore. Like cancer chemotherapy, the treatments for diabetes continues over a long period of time and has been the area of intensive protocol development. Unlike cancer chemotherapy, the treatment plan must handle multiple treatments in one day and deemphasizes the use of multiple drugs (although there are a variety of types of insulin). During 1987, using the medical literature and several internists in the medical computer science research group (Mark Frisse, Mark Musen, and Michael Kahn), we performed knowledge acquisition experiments for insulin treatment of diabetes. The proposed structure for the knowledge base was implemented using the object-oriented programming language upon which ONCOCIN has been based. These experiments, like those of adding more protocols to ONCOCIN, demonstrated the need for changes in the way that the knowledge base can access the time-oriented data base that records patient data and previous conclusions. The relationships between the different doses and types of insulin treatments will also require alternative ways of building treatment hierarchies. Thus, our initial experiments have shown that many of the elements of the ONCOCIN design are sufficiently general for other application areas, but that some specific elements (particularly the representation of temporal events) will have to be generalized. A description of our revised temporal representations will appear in a forthcoming thesis by Michael Kahn of University of California at San Francisco. During the coming year, we will continue our knowledge acquisition experiments and design a version of the E-ONCOCIN system that is separate from the ongoing "clinic version."

4 - OPAL: Graphical Knowledge Acquisition Interface

OPAL is a graphical environment for use by an oncologist who wishes to enter a new chemotherapy protocol for use by ONCOCIN or to edit an existing protocol. Although the system is designed for use by oncologists who have been trained in its use, it does not require an understanding of the internal representations or reasoning strategies used by ONCOCIN. The system may be used in two interactive modes, depending on the type of knowledge to be entered. The first permits the entry of a graphical description of the overall flow of the therapy process. The oncologist manipulates boxes on the screen that stand for various steps in the protocol. The resulting diagram is then translated by OPAL into computer code for use by ONCOCIN. Thus, by drawing a flow chart that describes the protocol schematically, the physician is effectively programming the computer to carry out the procedure appropriately when ONCOCIN is later used to guide the management of a patient enrolled in that protocol.

OPAL's second interactive mode permits the oncologist to describe the details of the individual events specified in the graphical description. For example, the rules for administering a given chemotherapy will vary greatly depending upon the patient's response to earlier doses, intercurrent illnesses and toxicities, hematologic status, etc. For example, one form permits the entry of an attenuation schedule for an agent based upon the patient's white count and platelet count at the time of treatment. Tables such as this are generally found in the written version of chemotherapy protocols. Thus OPAL permits oncologists to enter information using familiar forms displayed on the computer's screen. The contents of such forms are subsequently translated into rules and other knowledge structures for use by ONCOCIN.

4.1 - Status of the OPAL System

The OPAL is one of the few graphical knowledge acquisition systems ever designed for expert systems. Even fewer are designed to be used as the main method for entering knowledge as opposed to a proof of concept implementation. We have pursued three directions in the development of the OPAL system, also in response to the large number of protocols entered through this system during the last year. The first direction is the modification of graphical forms needed to allow the entry of facts that did not show up in the protocols used to test the initial version of OPAL. OPAL continues to assume that most of the knowledge to be entered will have very stereotyped forms, e.g., dose attenuations for most treatment toxicities are based on a comparison of only one laboratory measurement at a time, such as using the BUN to adjust for renal toxicity. We sometimes need much more complex ways of stating the scenarios in which dose adjustments may be necessary. This need has led us in a second direction, towards a "lower-level" rule entry approaching the syntax of the reasoning component of ONCOCIN, but using graphical input devices where applicable. A major accomplishment of this last year was to experimentally combine the OPAL and ONCOCIN programs into one working program, and to completely enter knowledge from OPAL using both the high level tools and lower level rule editors, but without needing to make changes at the ONCOCIN side of the system. Work on graphical replacements for low level rule entry is the master's project of Eric Sherman.

The OPAL program maps the information provided on the graphical forms into a complex data structure (called the IDS) that represents the required knowledge to specify the contents of a protocol. This data structure is used for copying information from one protocol to another, and as the basis for the creation of the ONCOCIN knowledge base. Our experiments with OPAL, and our intention to generalize OPAL use outside of oncology protocols, suggests that we reorganize the OPAL program to use a relational database to store its knowledge. We have patterned the database after an existing database query syntax. Because no databases exist for the InterLisp language upon which OPAL is based, we reimplemented the database from its written description. The database structure is completed, and was the basis for the PROTEGE knowledge acquisition experiments.

With changes in the future of dedicated lisp processors, we began to explore alternative platforms for developing OPAL-like systems. We have begun experiments using HyperCard on the MAC II. We also tested the usefulness of a "generic protocol viewing tool" based on a relational-database storage mechanism, in preparation for a OPAL design based on relational DBMS technology. We also performed several experiments aimed at developing a single underlying storage mechanism (termed a "cell") from which various interface systems including the Interviewer flowsheet, a generalized spreadsheet utility, and the OPAL schema entry flowchart interface could be derived.

5 - Generalized Knowledge Acquisition through PROTEGE

Mark Musen designed and implemented the first version of the PROTEGE knowledge-acquisition-system development tool. PROTEGE is used to collect information which describes the concepts (both entities and their relationships) in an application area for which a skeletal-planning type of expert system would be useful, concentrating on clinical trials. The system acquires the "ontology" of a domain through a series of fill-in-the-blank forms and a "flowchart-entry" tool. These concepts are then mapped onto a set of generic forms, which, in turn, create a knowledge acquisition tool for the application area.

PROTEGE makes use of the forms management system built for the original OPAL,

and a newly developed relational database management system written for the Xerox InterLisp-D workstations. The output of PROTEGE is an OPAL-like set of forms tailored to the special structures of the application area. To test these ideas, we first reimplemented portions of the OPAL interface from a high level description of oncology. After the translation process to the ONCOCIN reasoning program, a consultation was run that matched the manually built system. This experiment was then repeated for the area of hypertension protocols for which ONCOCIN had never been specifically designed. With some minor generalizations to the oncocin reasoner and interviewer, we were able to run a hypertension consultation.

6 - Speech Input to Expert Systems

6.1 - Prototype Speech Hardware/Software System

In 1987 we began a project to explore the integration of speech-recognition technology into the interface to ONCOCIN running on the XEROX lisp workstations. The project uses a commercially available continuous-speech-recognition product loaned by the vendor, Speech Systems, Inc. (SSI) of Tarzana, California. The speech recognizer consists of a custom processor, called the Phonetic Engine(R) and a suite of software modules called the Phonetic Decoder(TM). The Phonetic Decoder(TM) is running on a SUN 3/75.

The development of this project requires significant experience in distributed computing since the phonetic device, initial parsing software and the ONCOCIN system all reside on different pieces of hardware. One of the early steps is to allow the Lisp machine to remotely control the parsing software on the SUN. We built an interpreter for communicating between the speech software library running on the SUN and the Xerox Lisp machine. This interpreter reads lisp-style function calls corresponding to the speech library routines and returns lisp-style results as remote procedure call (RPC) mechanism. We then wrote the library of corresponding Lisp stub routines and a function to connect to the SUN workstation (through a programmatic TELNET) and start the server. In normal operation, we call the 'C'-based speech library routines from Lisp as if they were Lisp functions. We did a number of updates to both components of the server program (on the SUN and Lisp machine) as various SSI updates came out, including a major revision when SSI doubled the number of library routines, greatly increasing the server's capabilities.

Eventually we plan to replace the programmatic TELNET and custom RPC language with a new version of the system that uses the SUN RPC language. This will include the Xerox Lisp interface to the SUN RPC package developed by SUMEX; we were not able to use this interface initially as it did not exist early enough and required a release of Xerox Lisp that supports Common Lisp which the ONCOCIN system was not yet using. This eventual change will allow more efficient communications between the machines, allowing us to move larger data structures much faster. This will also tie in to the proposed method of communications between Lisp and non-Lisp routines that Xerox plans to use when their system migrates to the SUN workstations. The Lisp routines that correspond to the various SSI library routines will not appear to change to the application programs so they should not require any modification when the underlying RPC mechanism is changed.

We created a prototype system that permits users to navigate the graphical interface and enter clinical data using speech. The system uses the location of the cursor on the screen to provide a context for choosing candidate grammars with which to attempt recognition of a user's utterance. The system dynamically re-orders the list of candidate recognition grammars based on the dialog history. Albeit with limitations on the legal grammars, it is now possible to carry on most of the ONCOCIN data

acquisition steps using speech alone or speech plus pointing with the mouse. In addition, some elements such as the neural toxicities can be entered as textual descriptions and automatically encoded as one the 1-4 point scale used on flowsheet forms.

In order to translate an utterance back into an action that can occur in the ONCOCIN interface, we need the ability to reparse the text string returned by the SSI equipment. The SSI equipment uses (potentially complex) syntax, built up of various classifications, to understand sentences but returns just the ASCII component of the actual sentence; you can not get it in terms of the original classifications in the syntax (which are generally semantically significant). To overcome this problem, we devised a whole new syntax format based on Lisp and our OZONE object language in which one devises a grammar from which the SSI syntax files are generated. Then when the ASCII string is returned, the original syntax object is used to parse the string into a parse tree that relates directly to the grammar definition. We can now process the returned information at a much higher level than was possible with the simple ASCII text. In addition, the semantic elements of the parse can be added to syntactic structures used to encode the sentence.

The latest additions to this part of the system include exhaustive and random sentence generation as Lisp data structures (previously only partially available just on the SUN workstation) as well as word list generation from the syntax objects. These features generate both in-core representations as well as file representations which will potentially free us from any dependency on SSI's equipment (versus another manufacture's speech input device).

We wrote graphics software that plots the sentence returned by the speech equipment along with the changes in amplitude, pitch, accuracy scores and other information. This helps us to understand how SSI's 'black box' operates, particularly in situations where it fails. In addition, we wrote programs to make it possible to return all the best syntactic matches from the speech equipment. The current SSI software can either return the one best answer or all the possible word mappings (unfortunately regardless of syntax). We believe the latter would be useful when the most likely parse is associated with a low reliability score. We could then look at all the highest scoring candidates and evaluate what they do and do not have in common for creating a clarifying query to the speaker. The candidate sentences are composed from the word mapping data which is immense and needs to be compressed, filtered and manipulated to be of any use (the parser/generator is used as a filter to remove the non-syntactic word mappings). Unfortunately, though it runs faster with each improvement, this doesn't look reasonable for real time use; currently a three word utterance takes about two to four minutes to process into a candidate list of half a dozen possible sentences. It gives us an indication of what would be useful data structures to process on the SUN (to improve speed and lower data transmission).

6.2 - Speech Experiments

We are performing experiments to (1) enhance the system's grammars with a wider range of phrases clinicians actually use when talking to a computer and (2) gain insights into clinicians' models of spoken interaction with advice systems so that we may ground our interface design in observed practice.

In order to assess how physicians would speak to a computer in an ideal situation without constraints or prior assumptions, we are conducting a series of experiments which simulate continuous-speech understanding by computers. The setting of these experiments includes a hidden computer operator simulating the output of ONCOCIN if it had the ability to understand the spoken input, as well as a video camera to

record both audio and visual clues. Typed responses from the operator are translated back as actions on the computer display as well as audio responses through the use of a speech synthesizer. It appears to the subject as if the computer is understanding and responding to their speech. The physicians use ONCOCIN in the same manner as it is used in the clinic when they see patients, but with the added capability of speech input. These experiments enable us to both build up a basic vocabulary for the speech system as well as examine subtle linguistic issues to guide future directions.

One component of these experiments was the use of speech synthesis as an output medium. Using an inexpensive board based on the General Instrument ASCII-to-phoneme and phoneme-to-voice chip set, we built a driver for the Xerox Lisp machines. This driver included an application software transparent misspelling dictionary facility to correct for inaccuracies in the speech board.

7 - Object Language Support for ONCOCIN Project

We released a new version of our object language at the start of this past year which has proven to be the most stable and powerful version to date. There have been a number of minor bug fixes and several feature additions over the course of the year but for the most part the system as required much less attention than in previous years. The number of new systems being built on it (like our speech work) continues to increase. Future planning for the system consists of determining whether or not it should be converted to Common Lisp, based on whether object systems available under Common Lisp are sufficient for our needs, and if we do convert it what it would look like if properly integrated with that language.

8 - Personnel

Samson Tu has been primarily responsible for the design of E-ONCOCIN, Michael Kahn has developed the temporal representations used by the system, Clifford Wulfman has been involved with extensions to the data entry interface and the extensions to the interface in order to add speech input. Samson and Cliff were responsible for extensions to their programs to support the PROTEGE effort. David Combs has been involved with the knowledge acquisition interface and provided major programming support for the PROTEGE effort. Janice Rohn has been involved with the entry of protocols, interaction with physicians using the system, documentation of the system, and execution of the speech experiments. Ellen Isaacs, a Ph.D. student in Psycholinguistics has helped to design the speech experiments. Christopher Lane has developed the object-oriented systems software upon which the entire ONCOCIN system is designed and the systems software and parsing programs used in the speech project.